

REMARKS

Reconsideration of the above-identified patent application in view of the amendments above and the remarks following is respectfully requested.

Claims 1-25 are in this case. Claims 1, 6-11, 13-15, 18-22, 24 and 25 have been rejected under § 102(e). Claims 2-5, 12, 16, 17 and 23 have been objected to. Claims 2, 11, 12, 16, 22 and 23 have been canceled. Dependent claims 3 and 17 have been amended. New claims 26-30 have been added.

The claims before the Examiner are directed toward a method of link-level flow control and a network communication apparatus that uses that method. To each logical link, from a plurality of logical links in a physical link between a transmitter and a receiver, is assigned a respective maximum limit of transmission credit that corresponds to space available in a dynamically allocable portion of a buffer in the receiver. The sum of all the maximum limits corresponds to a space substantially larger than the dynamically allocable portion of the buffer. In response to traffic from the transmitter to the receiver, the receiver allocates credit to a logical link up to the logical link's respective maximum limit, as long as the total allocated credit does not exceed the dynamically allocable portion of the buffer. The transmitter transmits to the receiver responsive to this allocation, for example by not transmitting data whose size exceeds the allocated credit.

§ 102(e) Rejections – Fan et al. '165

The Examiner has rejected claims 1, 6-11, 13-15, 18-22, 24 and 25 under § 102(e) as being anticipated by Fan et al., US Patent No. 6,324,165 (henceforth, Fan et al. '165"). The Examiner's rejection is respectfully traversed.

Claims 11 and 22 have been canceled, thereby rendering moot the Examiner's rejection of these claims.

Fan et al. '165 teach a method of dynamically allocating bandwidth in a switch of an ATM network. It is clear that Fan et al. '165 do not anticipate the present invention, which is directed towards the dynamic allocation of buffer space in a network receiver

Furthermore, the present invention is not obvious from Fan et al. '165. The innovative concept of the present invention is that a receive buffer of a physical link in a network fabric such as an InfiniBand fabric need not be large enough to accommodate the total of all the credits that could be allocated to all the logical links of the physical link, but instead may be treated as a shared finite resource of the logical links. This innovative aspect of the present invention is expressed as follows in the specification of the above-identified patent application, on page 2 lines 16-31:

The credit limits are calculated by the receiver in such a manner as to guarantee that there will always be room available in the receive buffer for any packets sent by the transmitter within the applicable credit limits. This means that the receiver must leave an adequate amount of vacant buffer space for every one of the VLs that it serves. In order to maximize utilization of the bandwidth available on the physical link, the receiver preferably updates the transmitter's credit limits continually (and sends the appropriate flow control packets to the transmitter, while the transmitter is transmitting data), rather than waiting until the transmitter has used up its credit limit. It would thus appear that very large buffers are required at all of the switch ports in order to maintain efficient, wire-speed communications throughout the fabric.

Because Fan et al. '165 allocate bandwidth, not buffer space, Fan et al. are oblivious to the possibility of managing buffer space as a finite resource with respect to logical links. In any case, the strategy used by Fan et al. '165 to allocate bandwidth is not the strategy used by the present invention to allocate buffer space. Fan et al. '165 guarantee each of their connections on a link a minimum bandwidth M_i , such that the sum of the M_i is less than or equal to the link capacity C . Then, whenever the link has spare capacity, that spare capacity is shared dynamically among the active

connections. By contrast, the strategy of the present invention is to assign each logical link a maximum portion of the buffer, such that the sum of the maximum portions is substantially greater than the dynamically allocable portion of the receive buffer. Then each logical link is given as much buffer space as it requests, up to its maximum, subject to the constraint that the sum of the buffer space allocated to the links may not exceed the total available buffer space. It is not obvious from Fan et al. '165 that it would make any sense to assign entities, that compete for a finite resource, maximum shares of the resource whose sum exceeds the resource.

With independent claims 1 and 15 allowable in their present form, it follows that claims 6-10, 13, 14, 18, 21, 24 and 25, that depend therefrom, also are allowable.

Although claims 6 and 18 are allowable merely by virtue of depending from claims 1 and 15, Applicant takes the liberty of pointing out an additional reason why claims 6 and 18 are allowable.. Claims 6 and 18 address the issue of what is done if the resource allocated (bandwidth in the case of Fan et al. '165, buffer space in the case of the present invention) is inadequate. Fan et al. '165 react to a shortfall of the allocated resource by discarding cells. See for example column 3 lines 43-48:

The DRC scheduler detects congestion at bottleneck points in the switch and alleviates the congestion in a controlled manner by moving cell queueing towards the input side of the switch, where cell discard mechanisms such as early packet discard (EPD) and partial packed discard (PPD) may be applied to individual class queues. (emphasis added)

By contrast, the present invention, as recited in dependent claims 6 and 18, does not discard data, but instead withholds transmission. Even if it were obvious to extend the bandwidth allocation methods of Fan et al. '165 to buffer space allocation, one ordinarily skilled in the art would learn from Fan et al. '165 to discard data when the allocated buffer space is inadequate, not to withhold transmission of the data under such circumstances.

Objections

The Examiner has objected to claims 2-5, 12, 16, 17 and 23 as being based on rejected base claims. The Examiner has noted that claims 2-5, 12, 16, 17 and 23 would be allowable if rewritten in independent form including all the limitations of the base claim and any intervening claim.

Claim 2 has been rewritten in independent form, as new claim 26. Correspondingly, claim 2 has been canceled and claim 3 has been amended to depend from claim 26.

Claim 12 has been rewritten in independent form, as new claim 27. Correspondingly, claims 11 and 12 have been canceled.

Claim 16 has been rewritten in independent form, as new claim 28. Correspondingly, claim 16 has been canceled and claim 17 has been amended to depend from claim 28.

Claim 23 has been rewritten in independent form, as new claim 29. Correspondingly, claims 22 and 23 have been canceled.

Other New Claims

New claim 30 has been added

The aspect of the present invention that is addressed by new claim 30 relates to the locations of the transmitters and the receivers. In the case of Fan et al. '165, as illustrated in Figure 1, the transmitters are output ports **OP1-OP16** and **MOP** and output modules **OM1-OM16**, and the receivers are input ports **IP1-IP16** and input modules **IM1-IM16**, all of which are part of the same network switch that is based on core switch module **10**. By contrast, in the present invention, the transmitter and the receiver are in separate entities of the network. The description of the present invention in the specification of the above-identified patent application is in terms of

communication between a receiver in a port 24 of a switch 22 and a transmitter in a separate entity 27, as illustrated in Figure 1. That the receiver and the transmitter are in separate network entities is stated explicitly, on page 15 lines 15-18:

Fig. 4 is a flow chart that schematically illustrates a response of a receiver (port 24) to a data packet received from a transmitter (entity 27) over link 29...

read in conjunction with page 10 lines 12-18:

This port 24 transmits and receives packets over a full-duplex physical link 29 with another network entity 27, which likewise has queues 26 and 28. Typically, entity 27 comprises a port belonging to another switch, but it may also comprise a network host or peripheral device adapter, or a network device of another type, as is known in the art. (emphasis added)

New claim 30 is claim 15 as filed, with the additional limitation that the transmitter and the receiver are in separate entities of the network. Support for the new limitation is found in the above citations from the specification.


Amendments to the Specification

Inadvertent typographical errors on page 18 lines 10 and 11 have been corrected.

No new matter has been added.

In view of the above amendments and remarks it is respectfully submitted that independent claims 1, 15 and 26-30, and hence dependent claims 3-10, 13, 14, 17-21, 24 and 25 are in condition for allowance. Prompt notice of allowance is respectfully and earnestly solicited.

Respectfully submitted,



Mark M. Friedman
Attorney for Applicant
Registration No. 33,883

Date: July 5, 2004